

Reconfigurable Manifolds

Sarah Thompson

Alan Mycroft

`{sarah.thompson,alan,mycroft}@cl.cam.ac.uk`

Computer Laboratory, University of Cambridge, William Gates Building, 15 JJ Thomson Ave.,
Cambridge CB3 0FD, UK

1 Introduction

Spacecraft design is, without doubt, one of the most challenging areas of modern engineering. In order to be viable, spacecraft must mass relatively little, whilst being capable of surviving the considerable G-forces and vibration of launch. In space, they must withstand extreme temperatures, hard vacuum and high levels of radiation, for several years without maintenance.

Conventionally, spacecraft wiring harnesses are built with architectures that are fixed at the time of manufacture. They must therefore be designed to endure the lifetime of the mission with a very high probability, though the conventionally necessary redundant duplication of signals has significant implications for mass. Given that launch costs are typically in excess of \$30,000 per kg, reducing the mass of a spacecraft's wiring harness, without compromising reliability, is highly desirable. As a motivating example, the network cabling in the International Space Station (ISS) is known to mass more than 10 metric tonnes.

Recent advances in MEMS-based switching [9] have made it possible to consider the construction of *reconfigurable manifolds* – essentially, wiring harnesses that behave like macroscopic FPGA routing networks. Redundant wiring can be shared between many signals, thereby significantly reducing the total amount of cable required. Reconfigurability has a significant further benefit, in that it also allows adaptation to mission requirements that change over time, whilst also significantly reducing design time.

In a recent initiative, the US Air Force has been moving toward a *responsive space* paradigm which aims to reduce the time from design concept to launch (currently several years) to less than one week [7]. Such a target is unlikely to be achievable with existing bespoke one-off design techniques; a parts bin driven, plug-and-play approach to satellite construction will become essential. It must be possible to choose a satellite chassis of a size appropriate to the task in terms of accommodating sufficient manoeuvring propellant as well as the necessary instrumentation payload, then bolt everything together and have the resulting satellite ‘just work.’

We present an algorithm that allows such a reconfigurable manifold to be automatically self-configured, then dynamically tested in-situ, such that signals are automatically rerouted around non-functioning wires and switches as soon as faults are detected. Break-before-make switching is used in order to achieve transparency from the point of view of subsystems that are interconnected by the manifold, whilst also making it possible to achieve near-100% testability.




Figure 1: A typical near-earth small satellite configuration

1.1 Physical satellite wiring architectures

Conventionally, satellites are constructed with fixed wiring architectures. Reliability must therefore be engineered-in through multiple redundancy – duplication or triplication (or more) of signal paths is common, which carries with it an attendant mass penalty. Typically, one of two kinds of wiring architecture are common. Fig. 2 shows a typical passive backplane with multiple subsystems, each slotting in to a rack on separate cards¹. Wiring harnesses, in the sense that they exist in cars and aircraft as bundles of physical cables, tend to be avoided where possible.

Another common approach is shown in Fig. 3, where a single motherboard has a number of daughter boards attached to it on standoffs. Normally (though not visible in the diagram) these daughter boards plug directly into connectors on the motherboard, again avoiding the need for cables.

Typically, card frames have passive backplanes, which do not normally contain active electronics beyond perhaps some simple power regulation or line termination. Motherboard approaches more commonly include active electronics on the main board itself, though this is not a prerequisite.

¹Note that the image is representational – actual satellite hardware differs in detail




Figure 2: Card frame with backplane

1.2 Logical satellite wiring architectures

At a logical, block diagram level, fixed architecture satellite wiring harnesses typically follow the structure shown in Fig. 4. All of the main subsystems are attached to a motherboard or backplane that provides most of the necessary interconnection infrastructure, with external devices plugging directly into the relevant subsystems. All required redundancy must be in place from the outset. Typically, satellites are one-off designs, so any design changes before launch require physical modifications – of course, such changes *after* launch are typically impossible. As a further consequence of this approach, subsystem re-use is relatively uncommon, requiring considerable effort in terms of design, validation and verification, of the order of several years from concept to launch.

2 Reconfigurable manifolds

The responsive space paradigm [7] implies the requirement to move away from fixed architectures and their consequential design and validation costs toward an autonomous, self-organising approach. In essence, a *reconfigurable manifold* is a self-organising, self-testing, self-repairing replacement for a fixed architecture wiring harness. Ideally, at a system level, a spacecraft adopting this approach should have an architecture similar to that shown in Fig. 5.

Ideally, all wiring should be routed by the manifold rather than connected directly to subsystems. From the point of view of rapid construction, this is ideal – a subsystem




Figure 3: Motherboard with attached daughter boards

such as a gyroscope, star tracker, sun tracker or antenna could be bolted to the spacecraft chassis anywhere that is physically convenient, with all of the necessary wiring being ‘discovered’ and automatically routed after power-up.

2.1 Signal types

Spacecraft wiring harnesses (reconfigurable or otherwise) must be able to carry a wide variety of signals, varying in terms of power, voltage and bandwidth, with similarly variable electrical considerations in terms of impedance, end-to-end resistance, etc. Typical signal types found in satellites, along with example applications are listed as follows²:

Power Normally a single +28V DC unregulated supply rail powers the entire spacecraft, with local step-down regulators providing lower voltage high quality supply rails to each subsystem. Where higher voltages are necessary, e.g. to drive cryocoolers for low background noise imaging sensors, this is normally achieved with local step-up switching DC-DC converters.

Heavy current analogue High current feeds to torquer bars, motor drives, solenoid power, explosive bolts, etc.

Low current, low speed analogue Analogue sensor feeds, thermocouples, rough sun tracker photocells, etc.

²This list is not exhaustive

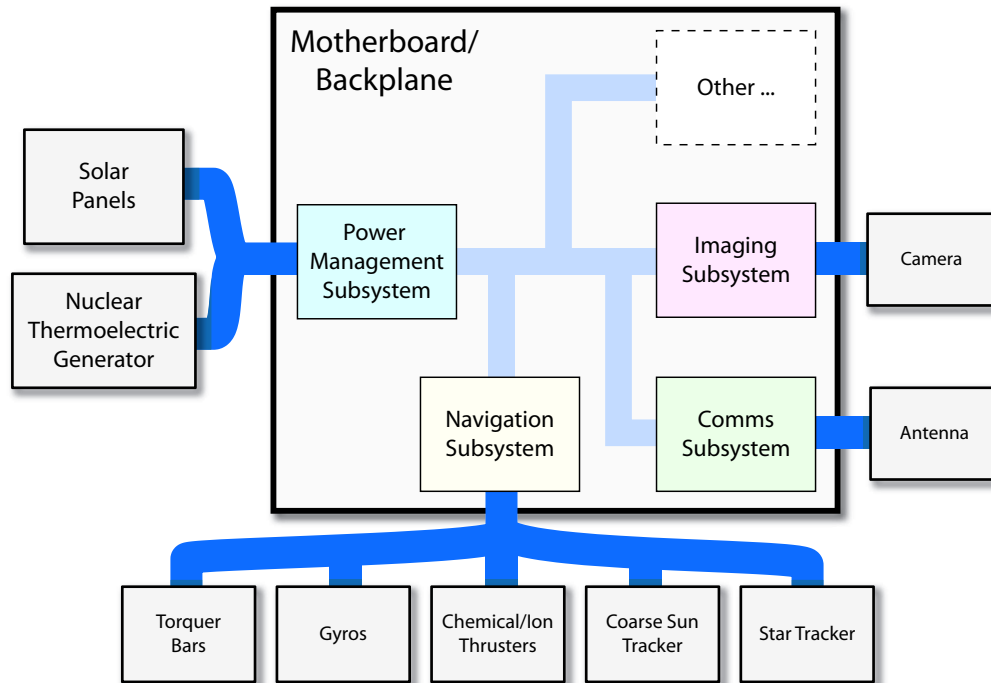


Figure 4: Conventional, fixed-architecture motherboard/backplane

Low current, high speed analogue Higher speed sensor wiring, video feeds from cameras and star trackers, etc.

Low speed digital Simple on/off telemetry sensors, e.g. mechanical limit switches.

High speed digital Digital communications between subsystems.

Low power microwave Radio receiver antenna feeds, low power radio transmitter antenna feeds.

High power microwave High power antenna feeds, ion thruster power cabling, etc.

Optical High speed network connectivity, lower speed sensor applications that require a significant degree of electrical isolation³.

No single switching architecture, at the time of writing, can accommodate more than a few of the above signal types.

2.2 Constructing practical reconfigurable manifolds

A practical reconfigurable manifold must encompass most, if not all, signal types in order to be effective. Since no single switch fabric is suitable, it makes sense to split the

³Optical switching is beyond the scope of this work and will not be discussed further

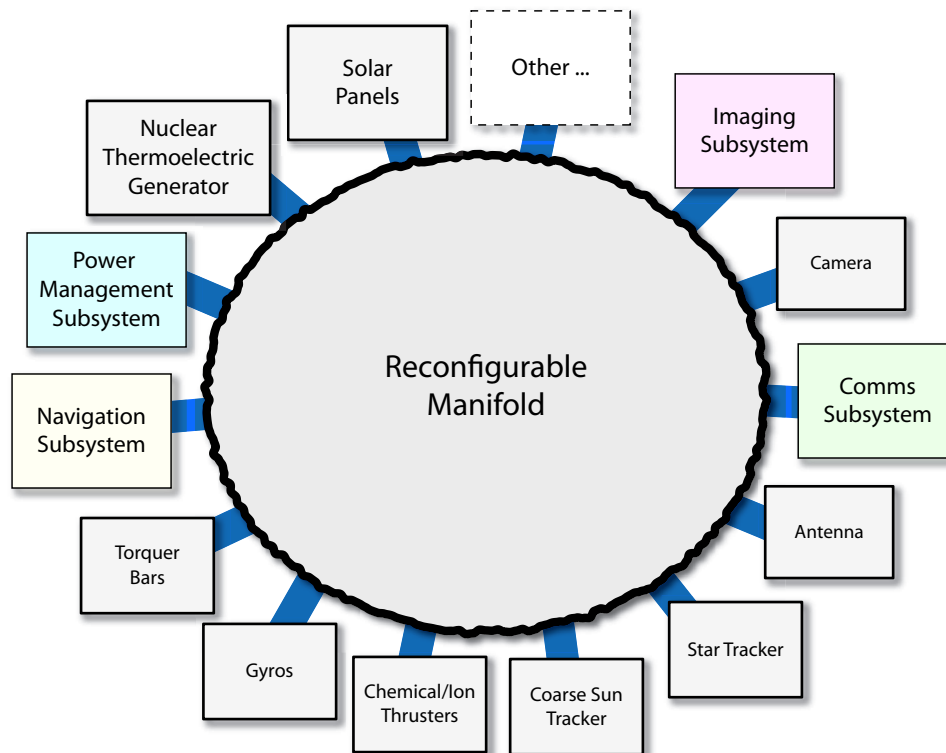


Figure 5: Reconfigurable manifold architecture

manifold into separate sub-manifolds, each of handling a different signal type, as shown in Fig. 6.

Some cross-connectivity between the sub-manifolds makes sense, since, for example, several MEMS relays could potentially be connected in parallel in order to switch heavier current, or DC-biased analogue routing with sufficient bandwidth could, in an emergency, be used to carry digital data.

Fig. 7 shows a reconfigurable manifold implemented as a replacement for a passive backplane or passive motherboard. In contrast with Fig. 4, external systems connect to the manifold rather than direct to the subsystems themselves. Configuring such a satellite might be as simple as installing cards in a backplane or motherboard in any convenient order, then plugging external devices into the manifold. Spare slots could, given sufficient mass budget, be used to provide extra redundancy simply by plugging in extra duplicate cards; appropriate firmware could potentially handle this automatically.

An alternative architecture is shown in Fig. 8. Rather than a single manifold routing between devices connected to its periphery, the manifold is itself distributed between the subsystems. Interconnection between subsystems is passive, with the subsystems cooperating to establish longer distance, multi-hop routes.

The single manifold approach is perhaps best suited to small satellites, whereas the (more complex, though more flexible and scalable) distributed approach lends itself to larger spacecraft such as large satellites, manned spacecraft, space stations or indeed

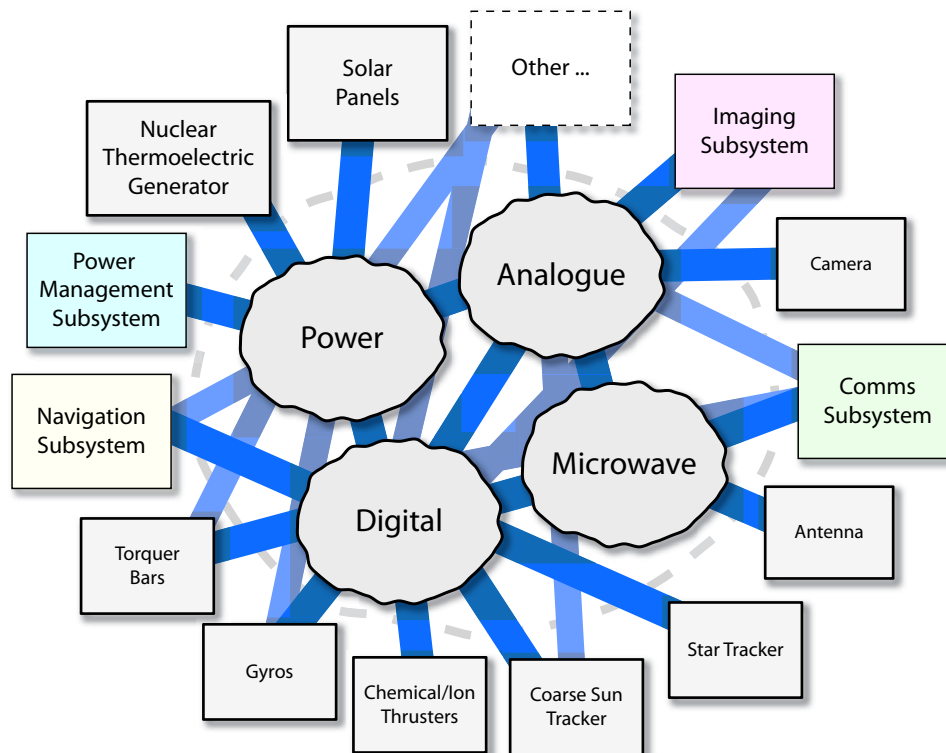


Figure 6: Separate routing networks for power, analogue, digital and microwave

also to terrestrial aircraft.

2.3 Switching technologies

Many switching technologies exist that differ considerably in capability:

FPGAs Field-programmable gate arrays can be used to route digital data, and are also comparatively cheap and readily available.

FPTAs Field-programmable transistor arrays [10] have some similarities to FPGAs, though they are aimed more closely at analogue applications. As with FPGAs, they are not intended from the outset as routing devices for use within a the switch fabric of a reconfigurable manifold, though it would seem feasible to apply them to the switching of low- to medium-speed analogue signals.

Digital Crossbar Switch ASICs A number of commercial, off-the-shelf (COTS) digital crossbar switch chips are available, though this application appears to be becoming dominated by FPGAs as a consequence of the larger FPGA manufacturers getting more directly involved by releasing support for using their devices in this way [2].

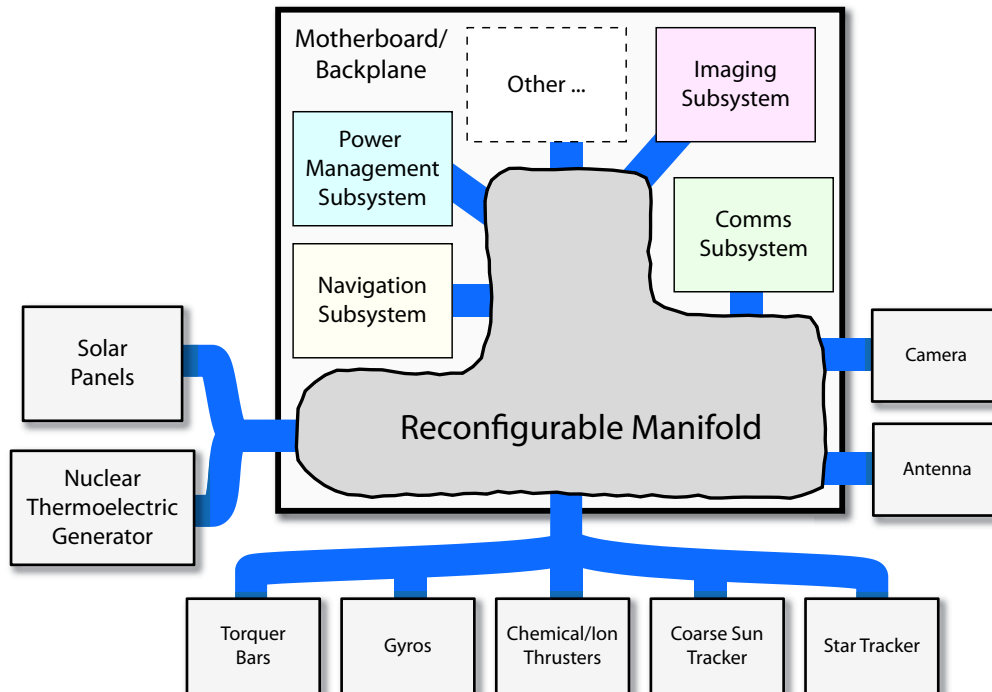


Figure 7: Reconfigurable manifold as a motherboard or backplane

Analogue Crossbar Switch ASICs Though not so widely supported as digital crossbar switch devices, analogue crossbar switches are available, mostly aimed at switching analogue video signals[1].

MEMS switches Micron-scale electromechanical switches have been demonstrated to be an effective candidate technology [9]. Though physically far larger than CMOS transistor-based switches, MEMS switches are nevertheless orders of magnitude smaller and lighter than full-size mechanical relays, and have excellent electrical characteristics that renders them capable of being applied to almost any low-current switching application, including microwave.

Electromechanical Relays Somewhat old-fashioned, relays are nevertheless capable of switching very heavy currents. They are sufficiently massive, however, that it is difficult to imagine them being used in large numbers in a spacecraft application.

Discrete MOSFET/IGBT Switching Large power transistors, both MOS and bipolar, are commonly used to switch heavy current and moderately high voltage (up to a few hundred volts and/or hundreds of amps) signals, particularly in motor drive applications. They exhibit high reliability and relatively good radiation hardness characteristics due to their very large (in comparison with ASICs) geometries, though their gate drive circuitry can be tricky to engineer. Though physically bulky, they nevertheless remain a useful possibility for constructing heavy current and/or power

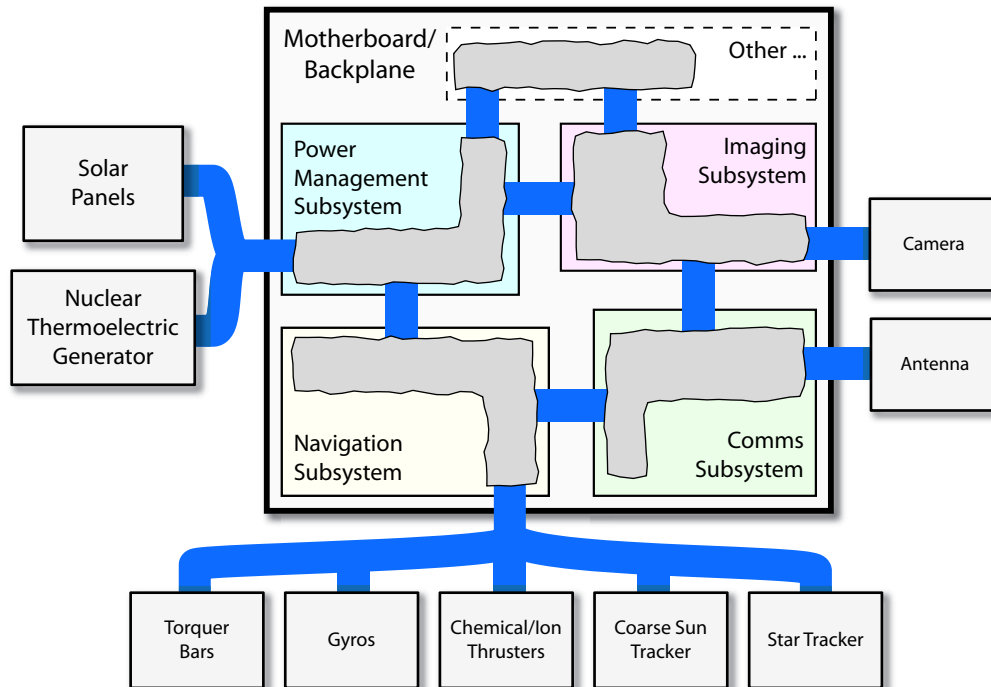


Figure 8: Reconfigurable manifold distributed across subsystems

switching networks.

Table 1 shows compatibility between switch technologies and signal types. The notation ‘?’ denoting ‘possibly compatible,’ indicates that, under normal circumstances, an automated routing algorithm would not normally attempt to make a connection of this type, though under certain circumstances, possibly only when authorised by a human, such connections might be made in the absence of more appropriate infrastructure. Normally, signals would be prioritised, so critical signals would almost always be routed, but less important connections may be degraded or even omitted. For example, a non-critical redundant temperature sensor might be disconnected in favour of keeping an instrument package running.

2.4 Routing architectures

The major alternative switching architectures that may be considered when designing a reconfigurable manifold are as follows:

Crossbar Switch An $M \times N$ grid of switches configured to provide a M -input, N -output routing network.

Permutation Network A *permutation network* performs an arbitrary permutation on N inputs, such that any possible reordering of the inputs is supported.

	FPGA	FPTA	Digital X-bar	Analogue X-bar	MEMS	Relays	MOSFET/IGBT
Power	×	×	×	×	?	✓	✓
Heavy current analogue	×	×	×	×	?	✓	✓
Low current, low speed analogue	×	✓	×	✓	✓	✓	?
Low current, high speed analogue	×	✓	×	✓	✓	?	?
Low speed digital	✓	✓	✓	✓	✓	✓	✓
High speed digital	✓	?	✓	?	✓	?	×
Low power microwave	×	×	×	×	✓	×	×
High power microwave	×	×	×	×	?	×	×

× – Not compatible ? – Possibly compatible ✓ – Compatible

Table 1: Compatibility between switch technologies and signal types

Ad-Hoc and Hybrid Approaches Practical considerations make it appropriate to consider the possibility of leveraging existing technologies, possibly in combination, to create reconfigurable manifolds. Though the result network topology and routing algorithms may be technically inferior to a purer design, economic considerations are nevertheless still important for practical designs.

Embedding into Networks of Arbitrary Topology Given a sufficiently large and complex graph, with nodes representing switches and edges representing wires, it is possible to compute a switch configuration that implements an arbitrary circuit.

Each approach is described in detail below.

2.4.1 Crossbar switches

Crossbar switches have a long history, having originally been introduced as a means of routing telephone calls through electromechanical telephone exchanges. Conceptually extremely simple, a crossbar switch is constructed from two sets of orthogonal wires (bus bars in telecommunications nomenclature), such that each crossing can be bridged by a switch. Fig. 9 depicts the circuit of a small 8×8 crossbar switch.

To route a particular input to a given output, all that is necessary is for the switch corresponding to that input and output to be closed. Crossbar switches are somewhat inefficient in terms of hardware requirements, and also in terms of providing more routing capability than is strictly necessary in many cases – it is possible, for example, to route a single input to any number of outputs, or to common inputs together. Achieving reliability is relatively straightforward, however – replacing each non redundant switch (Fig. 10) with a partially- or fully-redundant alternative (Fig. 11 or Fig. 12 respectively) allows single point failures to be recovered. A fully redundant switch configuration allows any of its four component switches to fail-open or fail-closed without affecting functionality. The partially redundant version only requires half as many switches, but is only safe against fail-closed faults – however, given one or more spare bus bars on each axis, fail-open faults can easily be patched around and are therefore still recoverable. In cost terms, building a fully-redundant $M \times N$ switch requires $4 \times M \times N$ switches, whereas the partially redundant approach requires $2 \times (M + 1) \times (N + 1)$ switches, though clearly the larger circuit is more fault-tolerant.




Figure 9: Crossbar Switch



Figure 10: Non-redundant switch



Figure 11: Partially redundant switch configuration




Figure 12: Fully redundant switch configuration

2.4.2 Permutation networks

Permutation networks are an alternative approach to routing that, in many cases, requires substantially fewer switches for a given number of inputs – rather than $O(N^2)$, they tend toward $O(N \log N)$, which can be a very significant advantage when the num-




Figure 13: 6-way permutation network

ber of inputs is large. Fig. 13 illustrates the concept with a 6-way permutation network. Its 15 switches can each be in either of two states: pass the inputs left to right unchanged, or swap them. For 6 inputs, a crossbar switch is likely to be cheaper, in that it is likely to require only 36 switches, in comparison with 60 for the permutation network shown in Fig. 13. However, for 1000 inputs, assuming $N \log_2 N$, approximately 40,000 switches are required, whereas a 1000×1000 crossbar switch would require 1 million switches.

Designing a permutation network can be somewhat baroque, though a useful relationship with *sorting networks* can be exploited. A sorting network is a sort algorithm that can be modified (if necessary) to allow its architecture to be predetermined, regardless of the data that it is given. Typically, a network is constructed whereby each *swap node* has two inputs and two outputs, where the outputs are swapped (if necessary) in order to respect a given partial order. Though the popular Quicksort is unsuitable, many other well-known sort algorithms, e.g. merge sort, bubble sort, transposition sort, bitonic sort or shell sort, can be adapted. Since a sort may also be seen as just a particular kind of permutation, sort networks – by definition – must be capable of performing permutations. Furthermore, since the data to be sorted might initially be in any order, a sort network must be capable of supporting all possible permutations – therefore, if a sort algorithm can be adapted to create a sort network of arbitrary dimension, it follows that an equivalently structured permutation network would also be capable of any possible permutation. Usefully, the underlying sort algorithm can be leveraged to efficiently generate switch configurations, as follows:

1. Let $\langle W, \sqsubset \rangle$ be a totally ordered set such that $|W|$ is the number of wires in the switch network, and each $w \in W$ represents exactly one input and one output.
2. Let the total bijective map $P : \wp(W \times W)$ represent the desired permutation to be implemented by the switch network.
3. Sort P with the underlying sort network, such that for each $(a, b) \in P$, a represents the input, and b represents the output. This can be achieved trivially by feeding tuples into the network ordered on a , then having the network sort these tuples ordered on b .
4. Note whether each swap node passed its data through unchanged, or whether it performed a swap. This gives the switch configuration for an isomorphic permutation network that performs an equivalent permutation.

Since suitable sort algorithms exist that have $O(N \log N)$ time complexity, computing a switch plan is therefore also an $O(N \log N)$ operation.

Permutation networks are nevertheless not guaranteed to be a better solution than crossbar switches, particularly when constructed as ASICs – their complex wiring reduces the effective advantage of their reduced switch count, particularly when considering that regular grids (crossbar switches being a particularly ideal example) are cheap and easy to lay out in comparison with the more spaghetti-like nature of large permutation networks. Limitations on chip packaging limit the number of wires that a single chip might be able to switch, and therefore also the number of switches that need sensibly be integrated in one die, reducing the impact of the $O(N^2)$ complexity problem with crossbar switches. However, when switches are large and/or expensive, as is the case with MEMS relays or any discrete component approach (e.g. full-size relays, MOSFETs, IGBTs), the reduction in component count could prove important.

2.4.3 Shuffle networks

Shuffle networks are essentially degenerate, incomplete permutation networks that do not support all possible permutations. They are perhaps best known in the parallel computing world, where they are commonly used as high speed inter-processor interconnect architectures. Omega networks, a commonly used shuffle network architecture, typically require some kind of blocking or queueing hardware at each swap node so that collisions can be arbitrated. Their incompleteness is probably not tolerable for our application, so they will not be considered further.

2.4.4 Ad-hoc COTS approaches

In some cases, COTS devices may be used to implement routing fabric. FPGAs, in particular, are ubiquitous, low cost and can be used (with appropriate considerations) in high radiation environments. There are a number of potential approaches:

1. Implement a general purpose crosspoint switch or permutation network as a HDL model, then synthesise it.
2. Generate HDL that routes the FPGA's inputs and outputs according to the desired switching plan, then synthesise the design.

The first option clearly limits the size of switch that can be implemented in a particular FPGA, though is inherently general purpose and can be reconfigured very rapidly. The second option is probably infeasible for embedded use at the time of writing due to the requirement for a complete tool chain in order to perform reconfiguration.

2.4.5 Embedding into networks of arbitrary topology

In this approach, a reconfigurable manifold is represented by a graph where its nodes represent switches and its edges represent wires. Embedding a desired circuit into such a network is essentially equivalent to computing a switch configuration. For the general case, this is a difficult computational problem that seems almost certainly to be in NP , with complexity rising exponentially with the number of switches in the network. Though this approach ultimately encompasses all others, in that both crossbar

switches and permutation networks may be seen as special cases, the difficulty of computing switching plans makes it unlikely that this approach could be feasible in practice.

2.5 Make-before-break switching

At the device level, make-before-break switching requires the capability to establish a new connection, in parallel, before an old connection is disconnected. Where a reconfigurable manifold is routing signals that should not be temporarily interrupted, make-before-break switching allows a connection to be moved to an alternative route transparently to the signal's endpoints.

Power, heavy current analogue, low-speed digital and low-speed analogue signals are all well suited to make-before-break switching, in that they are not particularly sensitive to minor changes in end-to-end resistance or discontinuities in impedance. However, high-speed digital, high-speed analogue, or (particularly) microwave signals need more careful consideration – in such cases, it may be necessary for the subsystems concerned to become involved in the routing process, at least from the point of view of being able to request that the manifold should not re-route particular signals during critical periods.

Crossbar switches support make-before-break switching by default: it is just necessary to turn on the switch for the new connection, waiting long enough (if necessary) for the switch to close fully and stop bouncing, then turn off the switch for the old connection. Implementing make-before-break switching in a permutation network is much more difficult, however, and will almost certainly require the network to be carefully designed (see Section 5.1).

In a reconfigurable manifold that does not alter its wiring plan after it has been initially configured, support for make-before-break switching is unnecessary – however, such a capability is essential in order to support continuous automated testing and fault recovery (see Section 4).

2.6 Grounding

Grounding of electronic systems within satellites is broadly similar to the grounding of Earth-based electronics; as-such, the same techniques and best practice applies in both cases. In satellites, grounding is particularly important because of the *charging effect*, whereby charged particles impacting the spacecraft impart a (potentially large) electric charge – careful grounding all conductive parts typically reduces or eliminates any consequential problems.

It is normal practice for a spacecraft to implement a ground network with a star topology – a single central grounding point is connected radially to the grounds on all subsystems. Cycles in the ground network are avoided, because they can form unwanted single-turn secondaries that may pick up hum or other unwanted noise from any heavy current subsystems in the vicinity.

Normally, grounds should not need to be switched by a reconfigurable manifold – a conventional, fixed, star ground topology should be sufficient for nearly all cases. Signals that are routed along shielded paths may require switchable ground lifts at one or both ends in order to avoid ground loops, though careful consideration of possible ground routing requirements may avoid this.

3 Self-organisation

In some circumstances, it is undesirable or even impossible to precalculate routing for a reconfigurable manifold. The responsive space paradigm requires that disparate subsystems should be able to be plugged together in any convenient manner, at which point they should self-organise and work together without human intervention. Achieving concept-to-launch times of the order of one week does not leave much time for anything other than physical assembly of the spacecraft, so the electronic subsystems must, of absolute necessity, not require a lengthy design process.

Self-organisation, at a fundamental level, requires subsystems to be able to discover each other, negotiate and configure any necessary wiring, and also to cooperate in maintaining the long-term reliability of the connectivity. These issues are discussed in detail in the remainder of this section.

3.1 ‘Space Velcro’

Some technologies absolutely require self-organisation in order to function at all. Fig. 14 is an electron micrograph of Joshi et. al.’s *Microcilia* concept [8, 11, 3]. MEMS technology is used to construct micron scale, articulated ‘cilia’ that are capable of manipulating small objects and of allowing the docking of small microsatellites. Assuming that electrical connections between the mated surfaces can be achieved, a self-organising, reconfigurable manifold based satellite could automatically configure any necessary connections during docking, then automatically recover the routing resources once the microsatellite has undocked.

Brei et. al. have investigated a passive interconnect architecture known as *Active Velcro* [6, 5, 4]. Fig. 15 illustrates the concept⁴. Mating, Velcro-like surfaces also contain a (possibly large) number of connectors, a proportion of which happen to make valid connections. Discovering these connections, then routing them via a reconfigurable manifold, potentially allows extremely straightforward ad-hoc construction. In manned spaceflight applications, an astronaut could connect or disconnect a piece of equipment simply by sticking or unsticking it to a Velcro-like pad⁵. In satellite applications, assuming that launch G force and vibration constraints are met, the same approach could allow extremely rapid construction and deployment.

3.2 Local routing

In a very small satellite, or within a single subsystem of a more complex satellite, routing may be exclusively *local*, i.e. switched only by a single level of switch networks. All connections in such a case would occur only to the edge of a single manifold, or cluster of sub-manifolds configured to act logically as a single manifold, with the consequence that the routing of all signals is equivalent only to routing across the manifold itself.

⁴Note that this is the author’s rendering, and is intended to be representational of the connectivity approach rather than an accurate physical description

⁵The use of Velcro to avoid small object floating around the cabin of manned spacecraft has long been standard practice.




Photo: John Suh, University of Washington

Figure 14: Microcilia Cell




Figure 15: Active Velcro

Computationally, routing for such an architecture is relatively trivial, with complexity of the order of $O(N^2)$ for a crossbar architecture or $O(N \log N)$ for a permutation network.

3.3 System level routing

Purely local routing requires a strict star architecture, with the manifold at the hub. This physical geometry does not suit all applications – in many cases, particularly in larger spacecraft, it is likely to be more appropriate to distribute the switching around the

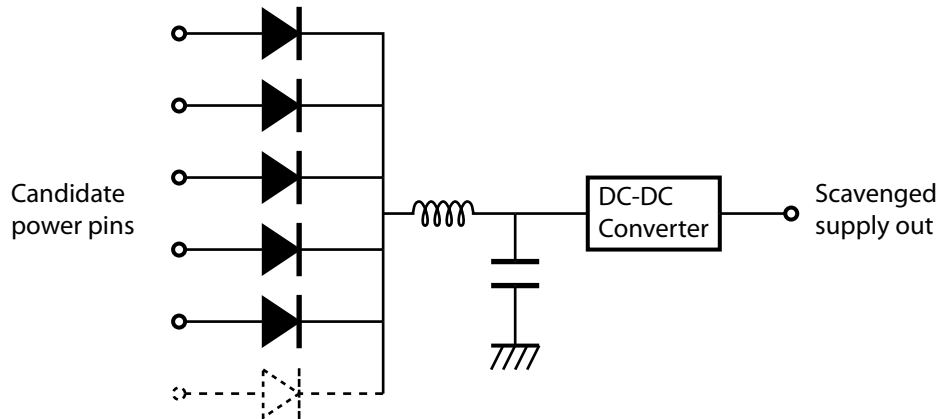


Figure 16: Power Scavenging Circuit

craft. Though it is theoretically possible to construct a large crossbar switch by ganging together smaller switches, this would be an expensive approach since the amount of inter-switch cabling would rise in proportion to the square of the number of switches. A more sensible and practical approach would be to construct a manifold-of-manifolds with an architecture resembling that of a circuit-switched telephone network – a number of manifolds handle primarily local connections internally, whilst handing off longer-distance connections via multicore trunk connections to other manifolds.

Computationally, the system level routing problem tends towards NP in the worst case (e.g. a manifold-of-manifolds where each manifold consists of exactly one switch and connectivity between manifolds is arbitrary is essentially the same problem that is discussed in Section 2.4.5), though the relatively small number of manifolds and relatively large amount of connectivity within each manifold is likely to minimise the consequences of this.

3.4 Dynamic discovery

The *dynamic discovery* of connections is something that is becoming increasingly common in general-purpose computing. The USB standard, for example, allows devices to be discovered and configured automatically without significant human intervention. From the point of view of reconfigurable manifolds, the dynamic discovery problem is somewhat trickier, in that it is necessary to first power up any neighbouring subsystems, establish contact with them (potentially with zero prior knowledge of their wiring configuration), negotiate any required connections, then route the necessary signals. As a second requirement, it is then necessary to continuously re-test the existing connectivity in order that faults can be corrected and that subsystems coming on line or going off line can be connected and disconnected correctly.

In this section, the requirements for achieving reliable dynamic discovery, continuous testing and fault recovery are discussed.

3.4.1 The chicken-and-egg problem

It is a truism that any automatic discovery algorithm can only possibly run on hardware that is itself powered up. However, if a subsystem's power connections have not yet been discovered and configured, it will not (yet) be powered up – hence there is a chicken-and-egg problem. Though no longer in common use, a well-known solution already exists. For many years, the most commonly used PC peripheral interface standards, RS232 and Centronics, both suffered from a design oversight – no power supply pins – that proved maddening for any hardware engineer attempting to design small peripherals without separate mains power supply connections. Designers nevertheless succeeded in working around the limitation by including circuits that scavenged power from the I/O pins themselves. The technique is illustrated in Fig. 16 – a diode network, effectively a large-scale generalisation of a full-wave rectifier circuit, synthesises power rails effectively by implementing a minimum/maximum function on the voltages that are present. The clamping, smoothing and DC-DC converter circuitry takes the potentially rather unpredictable raw output from the diode network and turns it into clean power that can be safely used to power up discovery circuitry prior to permanent routes being put in place.

Given suitable power scavenging circuits, a feasible power-up procedure for a large, manifold-of-manifolds architecture might be follows:

1. Power is applied to the first manifold through any arbitrary power pin.
2. The power scavenger circuit synthesises a suitable voltage rail for the embedded processor and discovery hardware responsible for the manifold.
3. All switches within the manifold are initialised to open circuit
4. The power connection is detected, then connected via the manifold, thereby disabling the diode network. This step avoids the inherent voltage drop across the diode network, whilst also reducing power consumption and heat dissipation slightly.
5. The manifold starts to listen for connection requests from other subsystems (see Section 3.4.3)
6. Power is temporarily routed to arbitrary pins on neighbouring subsystems that currently do not appear to be active, giving them the chance to power up and begin their own discovery process. They may request that power is supplied through a different pin, if necessary, or request that the existing pin should remain connected indefinitely⁶.

Eventually, all subsystems will be powered up, with the discovery process continuing to bring online all other necessary connections.

3.4.2 Watchdogs

It is standard practice for embedded processors in high reliability, mission critical and safety critical systems to be equipped with *watchdog circuits*, see Fig. 17.

⁶though it may be subject to change as part of the self-test algorithm

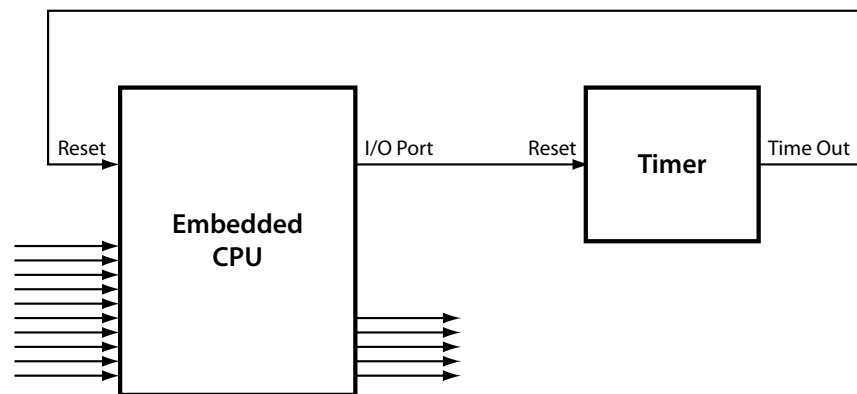


Figure 17: Typical watchdog circuit

A watchdog circuit is essentially a simple timer that is periodically reset by the host processor in such a way that, if the host processor happens to fail to reset it within a predetermined interval, the watchdog timer performs a hard reset on the host processor. Generally, this is integrated into a critical loop within the embedded software, such that if the program crashes this will cause the timer to fail to be reset, causing an automatic restart of the processor.

At a simplistic level, there is no reason why such a restart should cause problems for a manifold-of-manifolds architecture, though careful attention must be given to the following issues:

1. In the event of a watchdog reset, all external connections must be torn down, just in case the crash was itself caused by a faulty connection or, for example, by a SEE affecting the manifold itself.
2. Any negotiation protocol must be able to cope, e.g. by implementing timeouts, with connections going down without any corresponding explicit notification.

3.4.3 Discovery probe circuits

Connection discovery depends upon an ability to safely probe connections to find out what neighbouring subsystem they are connected to. The outline circuit shown in Fig. 18 shows how a suitable ‘discovery probe’ might be implemented. The circuit shows a UART (bidirectional serial interface) connected to a host processor, whose serial I/O ports (marked TxD and RxD) assume good quality, logic-level signals. On the transmit side, the signal is first buffered in order to protect the UART, then high pass filtered to achieve AC coupling and connected to the probe output via a resistor, whose value should be carefully selected in order to limit worst case current in the event of an accidental connection to a power or high current analogue signal to a level that can not cause damage. On the receive side, a similar current limiting resistor and high pass network protects the active components from direct connection to otherwise potentially damaging signals. A DC-coupled linear amplifier boosts the signal, then a Schmidt trigger (comparator with hysteresis) squares up the signal and raises it to logic levels suitable for the RxD input of

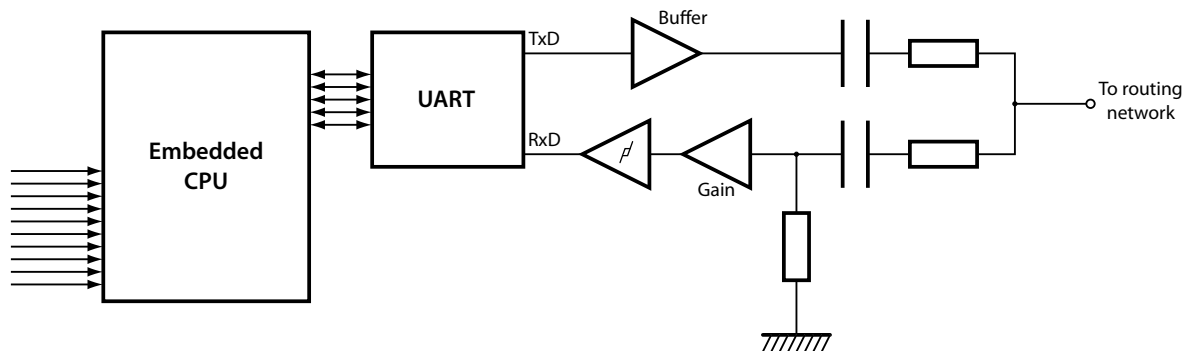


Figure 18: A possible discovery probe circuit




Figure 19: Typical packet format

the UART. Current limiting resistors should be chosen with values that are not too over-specified, since lower values are likely to result in better noise performance and higher achievable data rates.

In essence, the probe circuit is a simplified, extremely robust variation of a shared bus CSMA/CD network interface, in the style of 10Base2 Ethernet. AC coupling and a relatively high series resistance minimises the chance of damage due to accidental connection to higher voltage signals, whilst the ability to send and receive digital data without needing to switch between transmit and receive modes makes implementing higher level protocols relatively straightforward.

Sending serial data across AC coupled connections requires careful design of the low-level line protocol. Sending, for example, a long string of ones will cause the voltage to decay back to a centre value over a period of time that is determined by the time constant of the high pass filter. Similarly, a data packet that consists predominantly of ones (or zeros) will tend to shift away from the most common value, causing an unwanted DC bias and consequential reduction in noise margins. Typically this is addressed by arranging for the data encoding to implicitly retain an equal number of 0s and 1s – a trivial, though inefficient, approach is to spread an 8 bit byte across 16 bits, where each input bit corresponds to an inverted and a non-inverted copy in the output word. More efficient encodings exist that spread 2 bytes across 24 bits.

3.4.4 Line protocol

The main function of a suitable line protocol is to allow the discovery of connections, then to allow routing negotiation for signals. Probe circuits will typically alternate between sending packets that announce the identity of the relevant wire and listening for incoming packets that identify the other side of the connection. A suitable packet format is likely to follow the pattern shown in Fig. 19. Initially, a synchronisation waveform begins the transmission, whose purpose is to overcome any DC bias, whilst allowing the

receiving UART time to lock on to the data. A packet header follows, identifying the kind of packet that is being sent, followed by the packet payload and finally a checksum.

3.4.5 Connection establishment

Connections are established as follows (assuming a single manifold):

1. Both endpoints announce their identity, and announce the identifier of the signal that they wish to connect to.
2. Manifold detects the announcements
3. Manifold replies to both end points to say that the connection is being established, then ceases to probe either connection
4. Manifold establishes the connection, within a predetermined maximum time interval
5. Both endpoints are now free to use the connection.

More complex manifold-of-manifolds architectures will require more complex negotiation and routing, though the necessary protocols are likely to remain similar.

3.4.6 Stale connection tear-down

In the event that a subsystem crashes, stale connections should be torn down after a known time-out interval. The discovery probe protocol should also allow a connection to be torn down more rapidly by announcing that a neighbouring connection is no longer in use. Assuming that a dynamic testing and fault recovery process will be continuously applied, there is no requirement for a 'keep alive' protocol to ensure that valid connections stay up (see also Section4).

4 Dynamic testing and fault recovery

The same probe architecture necessary for discovery is also well suited to end-to-end testing of connections – if a connection is faulty (e.g. open circuit, shorted to ground or shorted to power), it will not be used, since the discovery process will fail to recognise it. As a consequence of this, at least for a short time after the discovery process has completed, all discovered connections may be regarded as functioning correctly. Over time, there is an increasing probability that, for example, permanent latch-up damage to a digital crossbar switch, may cause one or more connections to fail. This limitation can be avoided by constantly re-testing connections, ideally such that no connection may be established for a period longer than the minimum necessary to achieve the desired level of reliability.

4.1 Fault recovery protocol

There is actually no specific requirement to implement a fault recovery protocol as-such; the ability to set up and tear down connections, with make-before-break capabilities, is sufficient. Each end-point manifold should implement the following procedure (discovery and initial establishment of connections is assumed to have happened already):

1. Choose a signal on a round-robin basis
2. Establish a second route to the same remote end-point through the discovery protocol, which has the side-effect of ensuring that end-to-end connectivity is currently valid.
3. Connect the signal to the newly established route, at both ends, whilst leaving the original connection in place
4. Tear down the original connection
5. Repeat.

Note that in larger systems, connections between manifolds must always provide sufficient spare connections to allow the discovery protocol to remain in operation at all times.

The stale connection timeout (see Section 3.4.6) should be longer than the worst-case time necessary to cycle through all connections.

When a connection fails, it will be repaired automatically the next time that the fault recovery procedure cycles through the relevant signal, because the failed route will no longer be detected, so it will naturally fall out of the pool of available connections.

4.2 Graceful degradation

In a situation where cumulative failures have exceeded the number of available connections, it is sensible to define a graceful degradation strategy in order to maximise the spacecraft's remaining functionality. A simple approach is to rank all signals in order of importance, with signals toward the end of the list simply being disconnected if insufficient connectivity is available, though more sophisticated approaches may allow greater levels of recovery:

Routing signals on a less-ideal sub-manifold Normally, for example, digital data would be routed through dedicated digital switch networks. In the event that insufficient digital switching capacity remains, it is potentially feasible to route signals through spare capacity in other switch networks, e.g. via MEMS switching that would normally be used for microwave signals or via high speed analogue routes.

Multiplexing Manifolds could potentially be equipped with multiplexing hardware, in order that multiple low speed signals could be routed through a single connection. Though this may degrade any signals carried in this way, it may still be preferable to disconnecting signals entirely.

Emergency backup routing As an extension to the multiplexing approach, in an emergency backup routes could be established by non-standard means, such as via low power local digital radio links.

5 Conclusions

At the time of writing, this project is at a relatively early stage; nevertheless, it is possible to determine the following advantages of reconfigurable manifolds over conventional fixed-architecture spacecraft wiring harnesses:

Cost Reduction Since a reconfigurable manifold doesn't need to be designed from-scratch for each satellite, considerable cost reductions in terms of initial design, validation and verification are likely.

Reduction in Time To Launch (Responsive Space) Reduced design effort has a direct effect in terms of calendar time, potentially helping reduce a design process that is conventionally measured in years to just weeks or even days.

Possibility for Re-purposing After Launch If a spacecraft is no longer required for its initial purpose, given a modular design, it is quite likely that it could be re-purposed after launch at very low cost. For example, an imaging satellite with excess communications bandwidth could, assuming it has enough fuel, be shifted to another orbit to act as a communications relay.

Disaster Recovery Now legendary, the recovery of Apollo 13 after an explosion that deprived the command module of all three of its fuel cells and its entire oxygen reserve, with all crew alive and unhurt [12], was a direct consequence of heroic efforts to jury-rig the lunar lander's oxygen systems in order to keep the crew alive. A conventional satellite has no astronauts with a kit of spare parts available to make repairs – typically, failures tend to be terminal. A reconfigurable manifold offers great potential for jury-rigging the craft, either from Earth or possibly autonomously, so as to allow it to continue with some or all of its mission.

Mass reduction By sharing redundant wiring capacity across all subsystems, the total amount of copper necessary is reduced considerably in comparison with modular-redundant conventional wiring. At approximately \$30,000 per Kg to low earth orbit, even small savings can have considerable consequences in terms of cost.

5.1 Future Work

We conjecture that, in general, make-before-break switching is not feasible for permutation network based switch fabrics; further theoretical work is necessary in order to confirm this assumption. Ideally, it is hoped that a (probably non-optimal) permutation network architecture might be possible that can cope with interruption-free reconfiguration, though it is not clear at the time of writing how this might be achieved.

Many, if not all of the prerequisites for the practical construction of satellites based upon reconfigurable manifold technology are well-established, so the problem is primarily one of systems integration rather than difficult original R&D. The next step we intend to take is to build a software simulation of a reconfigurable manifold in order to test the feasibility of the approach. Beyond that, given appropriate funding and the necessary political will, it just remains to design a practical implementation and, hopefully, to trial it in space.

Acknowledgements

This work was supported by the US Air Force Office of Scientific Research Space Vehicles Directorate, through an EOARD grant. The first author wishes to thank AFOSR at Kirtland AFB, and Jim Lyke in particular, for their help and advice, without which this work would not have been possible.

References

- [1] *AD8116 - 200 MHz, 16 × 16 Buffered Video Crosspoint Switch*. Analog Devices, 2006. <http://www.analog.com/en/prod/0,2877,768>
- [2] *High Performance Crossbar Switch for Virtex-II and Virtex-II Pro FPGAs*. Xilinx, 2006. www.xilinx.com/esp/xbarswitch.htm.
- [3] BOHRINGER, K. F. A docking system for microsatellites based on microelectromechanical system actuator arrays. Tech. Rep. AFRL-VS-TR-2000-1099, US Air Force Research Laboratory, Space Vehicles Directorate, September 2000.
- [4] BREI, D., AND CLEMENT, J. Proof-of-concept investigation of active velcro for smart attachment mechanisms. Tech. Rep. AFRL-VS-TR-2000-1097, US Air Force Research Laboratory, Space Vehicles Directorate, September 2000.
- [5] BREI, D., AND CLEMENT, J. Velcro for smart attachment mechanisms. Tech. Rep. AFRL-VS-TR-2001-1104, US Air Force Research Laboratory, Space Vehicles Directorate, August 2001.
- [6] CLEMENT, J. W., AND BREI, D. E. Proof-of-concept investigation of Active Velcro for smart attachment mechanisms. In *In Proc. 42nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference and Exhibit* (2001). AIAA Paper 2001-1503 (AIAA Accession number 25238).
- [7] FOUST, J. Smallsats and standardization. *The Space Review* (2005).
- [8] JOSHI, P. B. On-orbit assembly of a universally interlocking modular spacecraft (7225-020). Tech. Rep. NASA SBIR 2003 Solicitation Proposal 03- II F5.03-8890, NASA, 2003.

-
- [9] LYKE, J., WILSON, W., AND CONTINO, P. MEMS-based reconfigurable manifold. In *Proc. MAPLD* (2005).
 - [10] STOICA, A., ARSLAN, T., KEYMEULEN, D., DUONG, V., GUO, X., ZEBULUM, R., FERGUSON, I., AND DAUD, T. Evolutionary recovery of electronic circuits from radiation induced faults. In *Proc. IEEE Conference on Evolutionary Computation* (2004), CEC.
 - [11] SUH, J. W., DARLING, R. B., BOHRINGER, K. F., DONALD, B., BALTES, H., AND KOVACS, G. T. A. SMOS integrated ciliary actuator array as a general-purpose micromanipulation tool for small objects. 1999.
 - [12] TURNILL, R. *The Moonlandings: an eyewitness account*. Cambridge, 2003.